

Cisco SDWAN Manager 3 Node Cluster Disaster Recovery

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Configure](#)

[Network Diagram](#)

[Configurations](#)

[Verification](#)

[How to verify replication leader node?](#)

[Validator \(vBond\) Password Update After Disaster Recovery Registration](#)

[Update Validator \(vBond\) Password](#)

[Adding New Validator \(vBond\) to Overlay After Disaster Recovery Registration](#)

[Upgrade Disaster Recovery Overlays](#)

[Before You Begin](#)

[Upgrade Process](#)

[Related Information](#)

Introduction

This document describes the stateful nature of Cisco vManage and its primary/secondary Designated Router (DR), enabling manual fail-over with auto data replication.

Prerequisites

Requirements

Cisco recommends that you have knowledge of vManage 3-node clusters.

Two separate vManage 3-node clusters must be configured and operational in order to proceed with disaster recovery. On the active cluster you must have validators and controllers onboarded. In case you have validator and controllers on the DR site, they must also be onboarded on the active cluster and not on the DR vManage cluster.

Cisco recommends that before registering disaster recovery, these requirements must be met:

- Ensure that the primary and the secondary node are reachable by HTTPS on a transport VPN (VPN 0).
- Ensure that Cisco vSmart Controllers and Cisco vBond Orchestrators on the secondary setup are connected to the primary setup.
- Ensure that the Cisco vManage primary node and secondary node are running the same Cisco

vManage version.

- Out-of-band cluster interface in VPN 0:
 - For each vManage instance within a cluster, a third interface (cluster link) is required besides the interfaces used for VPN 0 (transport) and VPN 512 (management).
 - This interface is used for communication and syncing between the vManage servers within the cluster.
 - This interface must be at least 1 Gbps and have a latency of 4ms or less. A 10 Gbps interface is recommended.
 - Both vManage nodes must be able to reach each other through this interface: be it a layer 2 segment or through layer 3 routing.
 - In each vManage, this interface must be configured in the GUI as a cluster interface(**Administration>Cluster Management**– indicate own out-of-band cluster interface IP address, user and password).
 - In order to allow Cisco vManage nodes to communicate with each other across data centers, enable TCP ports 8443 and 830 on your data center firewalls.
- Ensure that all services (application-server, configuration-db, messaging server, coordination server, and statistics-db) are enabled on both Cisco vManage nodes.
- Distribute all controllers, including Cisco vBond Orchestrators, across both primary and secondary data centers. Ensure that these controllers are reachable by Cisco vManage nodes that are distributed across these data centers. The controllers connect only to the primary Cisco vManage node.
- Ensure that no other operations are in process in the active (primary) and the standby (secondary) Cisco vManage node. For example, ensure that no servers are in the process of upgrading or attaching templates to devices.
- Disable the Cisco vManage HTTP/HTTPS proxy server if it is enabled. See [HTTP/HTTPS Proxy Server for Cisco vManage Communication with External Servers](#). If you do not disable the proxy server, Cisco vManage attempts to establish disaster recovery communication through the proxy IP address, even if Cisco vManage out-of-band cluster IP addresses are directly reachable. You can re-enable the Cisco vManage HTTP/HTTPS proxy server after disaster recovery registration completes.
- Before you start the disaster recovery registration process, navigate to the **Tools > Rediscover Network** window on the primary Cisco vManage node and rediscover the Cisco vBond Orchestrators.

Components Used

The information in this document is based on these software versions:

- Manager: 20.12.5
- Validator: 20.12.5
- Controller: 20.12.5
- cEdge: 17.12.5

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.


Background Information

Disaster recovery provides an administrator-triggered failover process. When disaster recovery is registered, data is replicated automatically between the primary and secondary Cisco vManage clusters. You manually perform a failover to the secondary cluster if needed.

Configure

Network Diagram

This figure illustrates the high-level architecture of the disaster recovery solution with a three node cluster.



Configurations

For more information on vManage Disaster Recovery, refer to [this](#) link.

The two separate 3-node-clusters are already created, assuming each SD-WAN manager has bare minimum configuration and certification part is completed.

```
vmanage2# show run system
system
 host-name          vmanage2
 system-ip          11.11.11.2
 site-id            1001
 admin-tech-on-failure
 no vrrp-advt-with-phymac
 sp-organization-name AAMIR-405707
 organization-name   AAMIR-405707
 upgrade-confirm     15
 vbond 10.105.60.104
```

```

vpn 0
interface eth0
ip address 10.105.60.102/24
ipv6 dhcp-client
tunnel-interface
no allow-service dhcp
allow-service dns
allow-service icmp
no allow-service sshd
no allow-service netconf
allow-service ntp
no allow-service stun
no allow-service https
!
no shutdown
!
interface eth1
ip address 89.89.89.2/24
no shutdown
!
ip route 0.0.0.0/0 10.105.60.1
!
vpn 512
interface eth2
ip address 10.105.60.192/24
no shutdown
!
ip route 0.0.0.0/0 10.105.60.1
!
vmanage2# show interface

```

VPN	INTERFACE	AF	IP ADDRESS	IF ADMIN STATUS	IF OPER STATUS	IF TRACKER STATUS	ENCAP TYPE	PORT	TYPE	MTU	HWADDR	SPEED MBPS	DUPLEX	TCP MSS ADJUST	UPTIME	RX PACKETS	TX PACKETS
0	eth0	ipv4	10.105.60.102/24	Up	Up	-	null	transport	-	00:0c:29:c0:37:03	1000	full	-	1:01:17:03	8806472	496731	
0	eth1	ipv4	89.89.89.2/24	Up	Up	-	null	service	-	00:0c:29:c0:37:0d	1000	full	-	1:01:16:59	16382852	15748084	
0	system	ipv4	11.11.11.2/32	Up	Up	-	null	loopback	-	-	1000	full	-	1:01:20:06	0	0	
0	docker0	ipv4	-	Down	Down	-	null	service	-	02:42:fb:fd:d4:86	1000	full	-	-	9	21	
0	cbr-vmanage	ipv4	-	Down	Up	-	-	-	-	02:42:c9:f5:28:c7	1000	full	-	-	-	-	
512	eth2	ipv4	10.105.60.192/24	Up	Up	-	null	mgmt	-	00:0c:29:c0:37:17	1000	full	-	1:01:16:59	994009	11814	

- Navigate to **Administration > Cluster Management** on both clusters and verify all nodes are in ready state.

DC vManage:

Cisco Catalyst SD-WAN

Select Resource Group

Administration - Cluster Management

Service Configuration

Service Reachability

Add Manager

Hostname	IP Address	Configure Status	Node Persona	UUID	
vmanage1	89.89.89.1	Ready	COMPUTE_AND_DATA	cb87a08e-079e-4394-81c3-e63c36ac22c0	...
vmanage2	89.89.89.2	Ready	COMPUTE_AND_DATA	8dc6c314-baca-40e7-a72c-94a3ebbe9d61	...
vmanage3	89.89.89.3	Ready	COMPUTE_AND_DATA	4a27ea41-3e1f-447c-baad-f6c3d07994d	...

DR-vManage:

Cisco Catalyst SD-WAN

Select Resource Group ▾

Administration · Cluster Management

≡

🔄

🔔

Service Configuration

Service Reachability

Add Manager

Hostname	IP Address	Configure Status	Node Persona	UUID	
DR-vmanage1	89.89.89.4	Ready	COMPUTE_AND_DATA	d78832e5-e6d3-4b6b-bf61-f923cf3c7282	...
DR-vmanage3	89.89.89.6	Ready	COMPUTE_AND_DATA	bf45f345-f2e-48ec-b8fd-0bb92427cc28	...
DR-vmanage2	89.89.89.5	Ready	COMPUTE_AND_DATA	c3e303a2-53d0-4525-901b-d96e9ce92875	...

- Navigate to **Administration>Disaster Recovery**. Click **Manage Disaster Recovery**.

Cisco Catalyst SD-WAN Administration - Disaster Recovery

Cluster Status

Active Cluster

Node	IP Address	Status
Disaster Recovery Not Configured		

Standby Cluster

Node	IP Address	Status
Disaster Recovery Not Configured		

Arbitrator

Node	IP Address	Status
Disaster Recovery Not Configured		

Details

Last Import:

Time to Import:

Size of Data:

Status:

History

Last Switch:

Reason for Switch:

Schedule

Replication Interval:

Switchover Threshold:

Manage Disaster Recovery Manage Password

Pause Disaster Recovery Pause Replication Delete Disaster Recovery

- In the pop-up window, fill the details for both primary and secondary vManage.

The IP addresses to be indicated are the out-of-band cluster interfaces IP addresses.

The credentials must be those of a netadmin user and they must not be changed once the DR is configured, unless it is deleted.

Manage Disaster Recovery

Connectivity Info Validator Info Recovery Mode Replication Schedule

Active Cluster

IP* 89.89.89.1

Username* admin

Password* ****

Standby Cluster

IP* 89.89.89.4

Username* admin

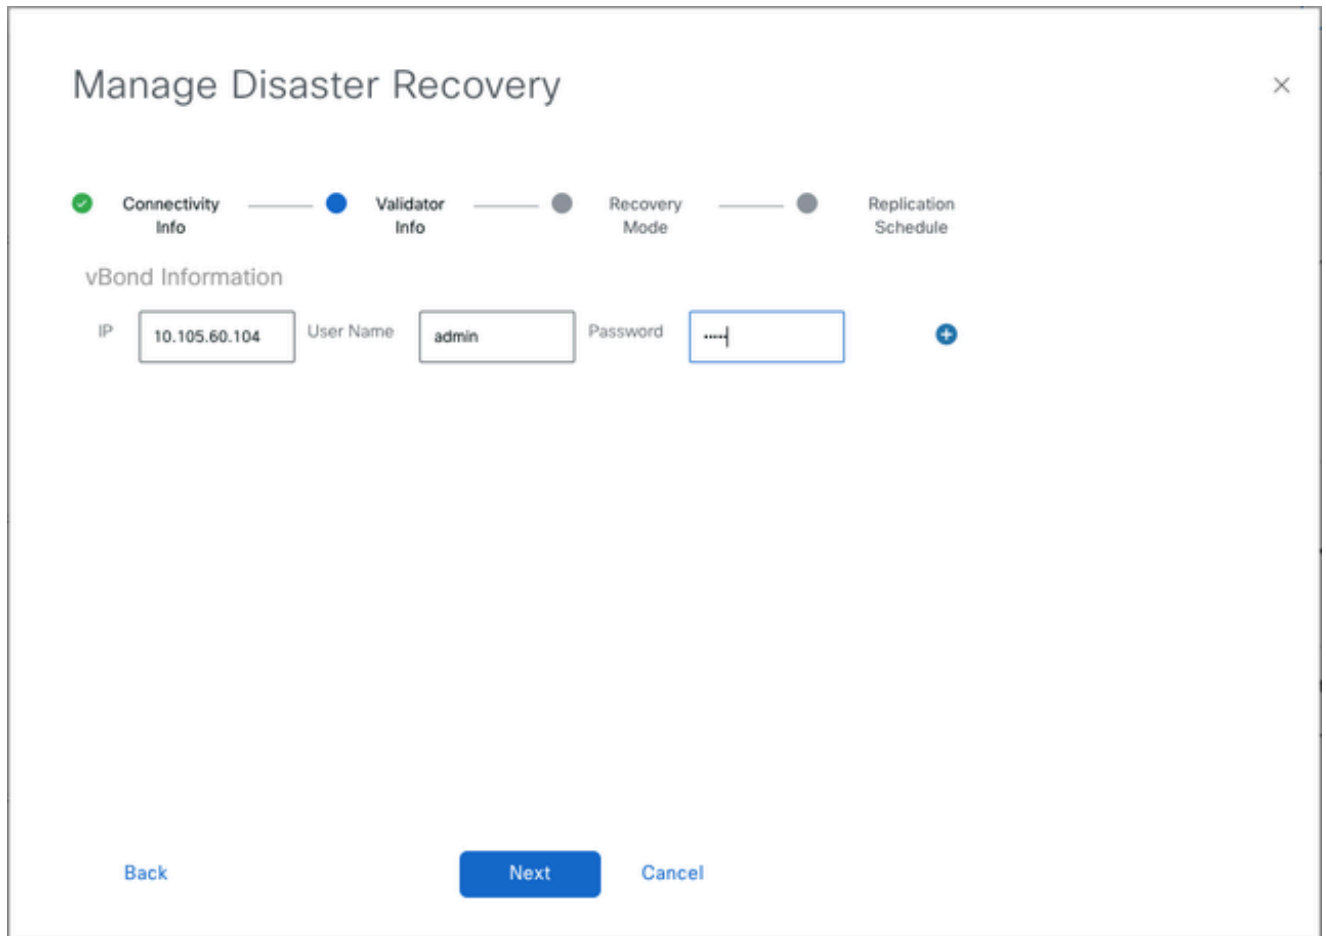
Password* ****

Next Cancel

Once filled, click **Next**.

- Fill the vBond controllers' details.

The vBond controllers must be reachable in the specified IP address via Netconf.



The image shows a 'Manage Disaster Recovery' dialog box with a close button (X) in the top right corner. At the top, there is a progress bar with four steps: 'Connectivity Info' (completed with a green checkmark), 'Validator Info' (current step with a blue dot), 'Recovery Mode' (grey dot), and 'Replication Schedule' (grey dot). Below the progress bar, the section is titled 'vBond Information'. It contains three input fields: 'IP' with the value '10.105.60.104', 'User Name' with the value 'admin', and 'Password' with masked characters '....'. To the right of the password field is a blue plus icon. At the bottom of the dialog, there are three buttons: 'Back' (disabled), 'Next' (active/blue), and 'Cancel' (disabled).

Once filled, click **Next**.

- In the Recovery Mode, choose **Manual**. The Automation mode is deprecated. Click **Next**.

Manage Disaster Recovery



Select Recovery Mode

- ☒ Manual ☐ Automation

Back

Next

Cancel

Manage Disaster Recovery

×

✓ Connectivity Info

✓ Validator Info

✓ Recovery Mode

● Replication Schedule

Start Time

12:00

AM

Replication Interval

15 mins

Back

Save

Cancel

Set the value and click **Save**.

- The DR Registration starts now. Click the refresh button to manually refresh the state and the progress logs. This process can take up to 20-30 minutes.

Cisco Catalyst SD-WAN
Select Resource Group
Administration - Disaster Recovery

Disaster Recovery Registration
Total Task: 1 | Success: 1
Device Group (1)
Search Table

Status	Chassis Number	Hostname	Message
Success	-	-	Data Centers Register


View Logs

[4-Jul-2025 4:38:41 UTC] [4-Jul-2025 4:33:03 UTC] Vmanage 89.89.89.5 has successfully restarted.
[4-Jul-2025 4:38:41 UTC] [4-Jul-2025 4:33:22 UTC] Restart initiated. Waiting for Vmanage 89.89.89.5 to come up.
[4-Jul-2025 4:38:41 UTC] [4-Jul-2025 4:36:03 UTC] Vmanage 89.89.89.5 has successfully restarted.
[4-Jul-2025 4:38:41 UTC] [4-Jul-2025 4:36:03 UTC] 2 vmanages have successfully registered and restarted. Restarting current vmanage 89.89.89.4
[4-Jul-2025 4:38:42 UTC] Restarting Primary DC
[4-Jul-2025 4:38:42 UTC] Restarting Local DataCenter
[4-Jul-2025 4:38:42 UTC] Restarting Vmanage 89.89.89.3
[4-Jul-2025 4:39:02 UTC] Restart initiated. Waiting for Vmanage 89.89.89.3 to come up.
[4-Jul-2025 4:40:13 UTC] Vmanage 89.89.89.3 has successfully restarted.
[4-Jul-2025 4:40:13 UTC] Restarting Vmanage 89.89.89.2
[4-Jul-2025 4:43:34 UTC] Restart initiated. Waiting for Vmanage 89.89.89.2 to come up.
[4-Jul-2025 4:52:38 UTC] Vmanage 89.89.89.2 has successfully restarted.
[4-Jul-2025 4:52:40 UTC] 2 vmanages have successfully registered and restarted. Restarting current vmanage 89.89.89.1

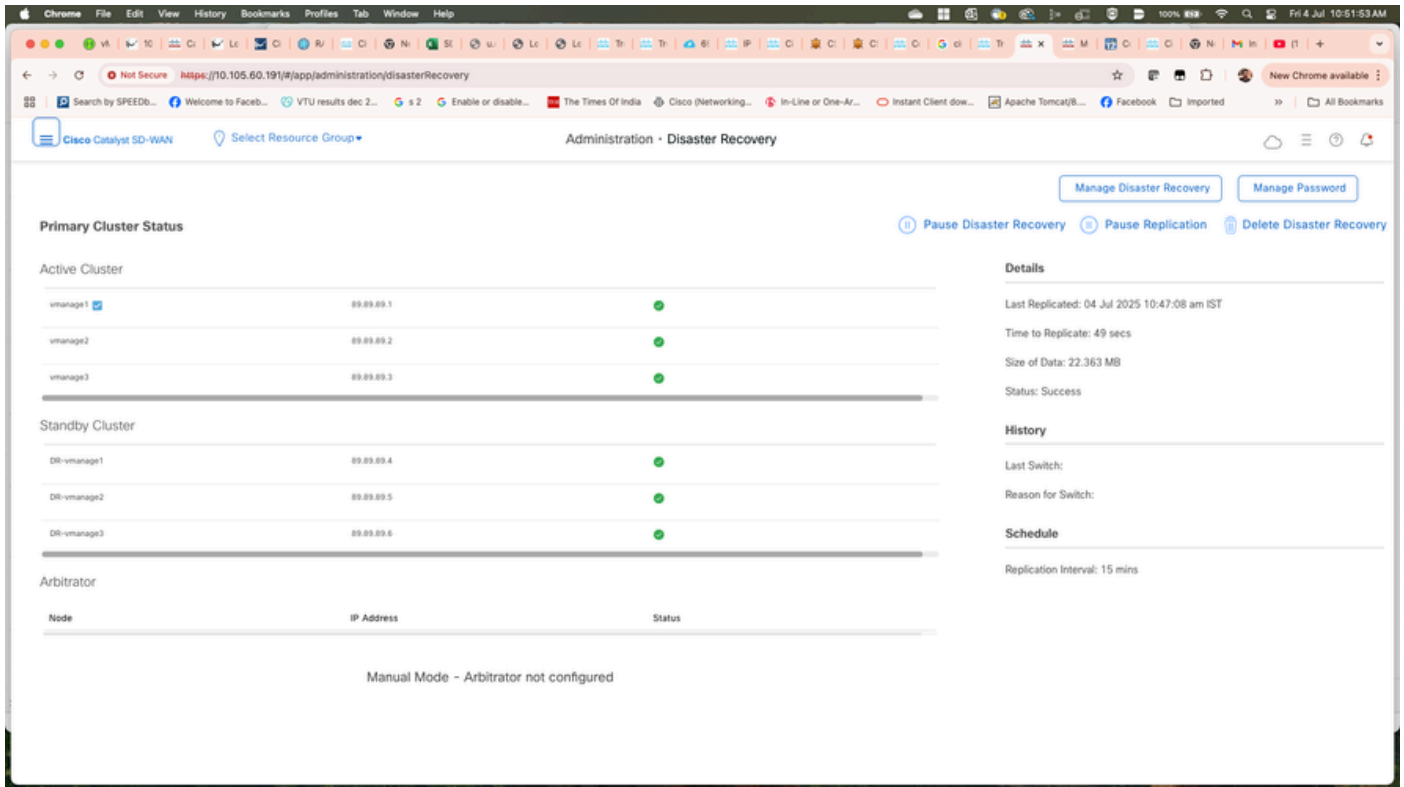
Close

Verification

- Navigate to **Administration>Disaster Recovery** in order to see the Disaster Recovery status and when the data was replicated last time.



Note: In this scenario, replication took only 49 seconds because the lab environment has a small database. However, replication can take several hours depending on the database size. Additionally, it can require a few cycles to achieve successful replication.



Verify disaster recovery log in both clusters.

DC-vmanage (9a15f979-d613-4d75-97bf-f7d4124bc687 is export ID)

```
vmanage1:/var/log/nms$ cat vmanage-disaster_recovery.log | grep 9a15f979-d613-4d75-97bf-f7d4124bc687
04-Jul-2025 05:17:08,297 UTC INFO [] [] [DataReplicationManager] (pool-232-thread-1) || Export ID Gener
04-Jul-2025 05:17:58,431 UTC INFO [] [] [DisasterRecoveryAlarmsDAO] (pool-232-thread-1) || AlarmsDAO::a
04-Jul-2025 05:17:58,722 UTC INFO [] [] [DataReplicationManager] (pool-232-thread-1) || Sending the imp
04-Jul-2025 05:17:59,081 UTC INFO [a17a50ae-e6d3-401c-9d34-7c9423a5dd5a] [vmanage1] [DisasterRecoveryRe
04-Jul-2025 05:21:06,515 UTC INFO [a456da19-9868-42e1-b3e7-9cb7ef3bdb81] [vmanage1] [DisasterRecoveryRe
vmanage1:/var/log/nms$
```

DR-Vmanage

```
DR-vmanage1:/var/log/nms$ cat vmanage-disaster_recovery.log | grep 9a15f979-d613-4d75-97bf-f7d4124bc687
04-Jul-2025 05:15:23,296 UTC INFO [] [] [DataReplicationManager] (Thread-366) || Payload received for d
04-Jul-2025 05:15:23,298 UTC INFO [] [] [DataReplicationManager] (Thread-366) || destinationURL dataser
04-Jul-2025 05:15:24,040 UTC INFO [] [] [DisasterRecoveryAlarmsDAO] (Thread-366) || AlarmsDAO::addAlarm
04-Jul-2025 05:15:24,170 UTC INFO [] [] [DataReplicationManager] (Thread-366) || Downloaded replication
04-Jul-2025 05:15:24,171 UTC INFO [] [] [DisasterRecoveryManager] (Thread-366) || Sending rpc message to
04-Jul-2025 05:15:24,216 UTC INFO [] [] [DisasterRecoveryManager] (Thread-366) || Sending message to de
04-Jul-2025 05:15:24,245 UTC INFO [] [] [DisasterRecoveryManager] (Thread-366) || Waiting for copyRepli
```

```
04-Jul-2025 05:18:19,545 UTC INFO [] [] [DataReplicationWorker] (Thread-366) || Successfully Deleted Im
04-Jul-2025 05:18:19,643 UTC INFO [] [] [DisasterRecoveryAlarmsDAO] (Thread-366) || AlarmsDAO::addAlarm
04-Jul-2025 05:18:19,707 UTC INFO [] [] [DataReplicationManager] (Thread-366) || Successfully imported
04-Jul-2025 05:18:19,716 UTC INFO [] [] [DisasterRecoveryManager] (Thread-366) || Sending rpc message to
04-Jul-2025 05:18:19,849 UTC INFO [] [] [DisasterRecoveryManager] (Thread-366) || Sending message to de
```

How to verify replication leader node?

- Use the next API in order to findout replication leader node on both clusters:

<https://<vmanage-ip>/data service/entity ownership/tree>.

For DC cluster:

Replication node is cb87a08e-079e-4394-81c3-e63c36ac22c0 which is node1, verify it from show control local-properties.

```
https://10.105.60.192/dataservice/entityownership/tree

{"bucket_5": [{"entityName": "notificationService", "bucket": "/bucket_4", "owner": "4a27ea41-3e1f-447c-baad-f6c3d07994d"}], "/bucket_3": [{"entityName": "TelemetryHealthCollectorSchedulerService", "bucket": "/bucket_3", "owner": "4a27ea41-3e1f-447c-baad-f6c3d07994d"}, {"entityName": "ContainerLManager", "bucket": "/bucket_3", "owner": "4a27ea41-3e1f-447c-baad-f6c3d07994d"}], "/bucket_2": [{"entityName": "ReportStatusMonitor", "bucket": "/bucket_2", "owner": "4a27ea41-3e1f-447c-baad-f6c3d07994d"}, {"entityName": "singleton-service-selector", "bucket": "/bucket_2", "owner": "4a27ea41-3e1f-447c-baad-f6c3d07994d"}, {"entityName": "correlation-engine-manager", "bucket": "/bucket_2", "owner": "4a27ea41-3e1f-447c-baad-f6c3d07994d"}], "/bucket_1": [{"entityName": "schedules-service-singleton-selector", "bucket": "/bucket_1", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}, {"entityName": "TelemetryDataCollectorSchedulerService", "bucket": "/bucket_1", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}], "/bucket_0": [{"entityName": "CRLSettingsManager", "bucket": "/bucket_0", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}, {"entityName": "dcaServiceOwner_default", "bucket": "/bucket_0", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}], "/bucket_9": [{"entityName": "elasticSearchIndexMigrationService", "bucket": "/bucket_9", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}, {"entityName": "drNodeStateService", "bucket": "/bucket_9", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}, {"entityName": "manageK5Service", "bucket": "/bucket_9", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}], "/bucket_8": [{"entityName": "smartLicenseIngenManagerService", "bucket": "/bucket_8", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}, {"entityName": "TelemetryConfigurationCollectorsSchedulerService", "bucket": "/bucket_8", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}, {"entityName": "TunneHealthTask", "bucket": "/bucket_8", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}], "/bucket_7": [{"entityName": "all", "bucket": "/bucket_7", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}, {"entityName": "mpaadapter", "bucket": "/bucket_7", "owner": "8dc6c314-baca-40e7-a72c-94a3ebbe9d61"}], "/bucket_COMPUTE_0": [{"entityName": "drReplicationService", "bucket": "/bucket_COMPUTE_0", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}], "/bucket_6": [{"entityName": "multiCloudScheduleManager", "bucket": "/bucket_6", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}, {"entityName": "CleanScheduleManager", "bucket": "/bucket_6", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}, {"entityName": "sdwaaasScheduleManager", "bucket": "/bucket_6", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}, {"entityName": "singleton-service-client", "bucket": "/bucket_6", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}], "/bucket_6": [{"entityName": "drReplicationService", "bucket": "/bucket_COMPUTE_0", "owner": "cb87a08e-079e-4394-81c3-e63c36ac22c0"}]}
```

Similarly for DR-vManage, replication node is d78832e5-e6d3-4b6b-bf61-f923cf3c7282.

```
https://10.105.60.195/dataservice/entityownership/tree

{"bucket_5": [{"entityName": "notificationService", "bucket": "/bucket_4", "owner": "bf45f345-f12e-48ec-b8fd-0bb92427cc28"}], "/bucket_3": [{"entityName": "TelemetryHealthCollectorSchedulerService", "bucket": "/bucket_3", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}, {"entityName": "ContainerLManager", "bucket": "/bucket_3", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}], "/bucket_2": [{"entityName": "ReportStatusMonitor", "bucket": "/bucket_2", "owner": "bf45f345-f12e-48ec-b8fd-0bb92427cc28"}, {"entityName": "singleton-service-selector", "bucket": "/bucket_2", "owner": "bf45f345-f12e-48ec-b8fd-0bb92427cc28"}, {"entityName": "correlation-engine-manager", "bucket": "/bucket_2", "owner": "bf45f345-f12e-48ec-b8fd-0bb92427cc28"}], "/bucket_1": [{"entityName": "schedules-service-singleton-selector", "bucket": "/bucket_1", "owner": "bf45f345-f12e-48ec-b8fd-0bb92427cc28"}, {"entityName": "TelemetryDataCollectorSchedulerService", "bucket": "/bucket_1", "owner": "bf45f345-f12e-48ec-b8fd-0bb92427cc28"}], "/bucket_0": [{"entityName": "CRLSettingsManager", "bucket": "/bucket_0", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}, {"entityName": "dcaServiceOwner_default", "bucket": "/bucket_0", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}], "/bucket_9": [{"entityName": "elasticSearchIndexMigrationService", "bucket": "/bucket_9", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}, {"entityName": "drNodeStateService", "bucket": "/bucket_9", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}, {"entityName": "manageK5Service", "bucket": "/bucket_9", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}], "/bucket_8": [{"entityName": "smartLicenseIngenManagerService", "bucket": "/bucket_8", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}, {"entityName": "TelemetryConfigurationCollectorsSchedulerService", "bucket": "/bucket_8", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}, {"entityName": "TunneHealthTask", "bucket": "/bucket_8", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}], "/bucket_7": [{"entityName": "all", "bucket": "/bucket_7", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}, {"entityName": "mpaadapter", "bucket": "/bucket_7", "owner": "c3e303a2-53d0-4525-901b-d96e9ce92875"}], "/bucket_6": [{"entityName": "multiCloudScheduleManager", "bucket": "/bucket_6", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}, {"entityName": "PerfMonTask", "bucket": "/bucket_6", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}, {"entityName": "CleanScheduleManager", "bucket": "/bucket_6", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}, {"entityName": "sdwaaasScheduleManager", "bucket": "/bucket_6", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}, {"entityName": "singleton-service-client", "bucket": "/bucket_6", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}], "/bucket_COMPUTE_0": [{"entityName": "drReplicationService", "bucket": "/bucket_COMPUTE_0", "owner": "d78832e5-e6d3-4b6b-bf61-f923cf3c7282"}]}
```

Validator (vBond) Password Update After Disaster Recovery Registration

If you change the vBond password after disaster recovery registration is complete, a switchover fails because the vBond password is not updated on the secondary cluster, which still retains the old vBond password.

[04-July-2025 6:47:35 UTC] Unshut control tunnel on the standby vManage.

[04-July-2025 6:47:36 UTC] Sleeping for 10 seconds to ensure control tunnel is fully up and functional

[04-July-2025 6:47:55 UTC] Failed to activate the cluster. Vbond is unreachable

=====

```
04-July-2025 06:47:55,206 UTC ERROR [89b008fa-2c1b-4f78-b093-ed1fa1f06b71] [vManage20-14-DR] [DisasterR
at com.viptela.vmanage.server.device.common.NetConfClient.connect(NetConfClient.java:255) ~[vmanage-ser
at com.viptela.vmanage.server.device.common.NetConfClient.<init>(NetConfClient.java:114) ~[vmanage-serv
```

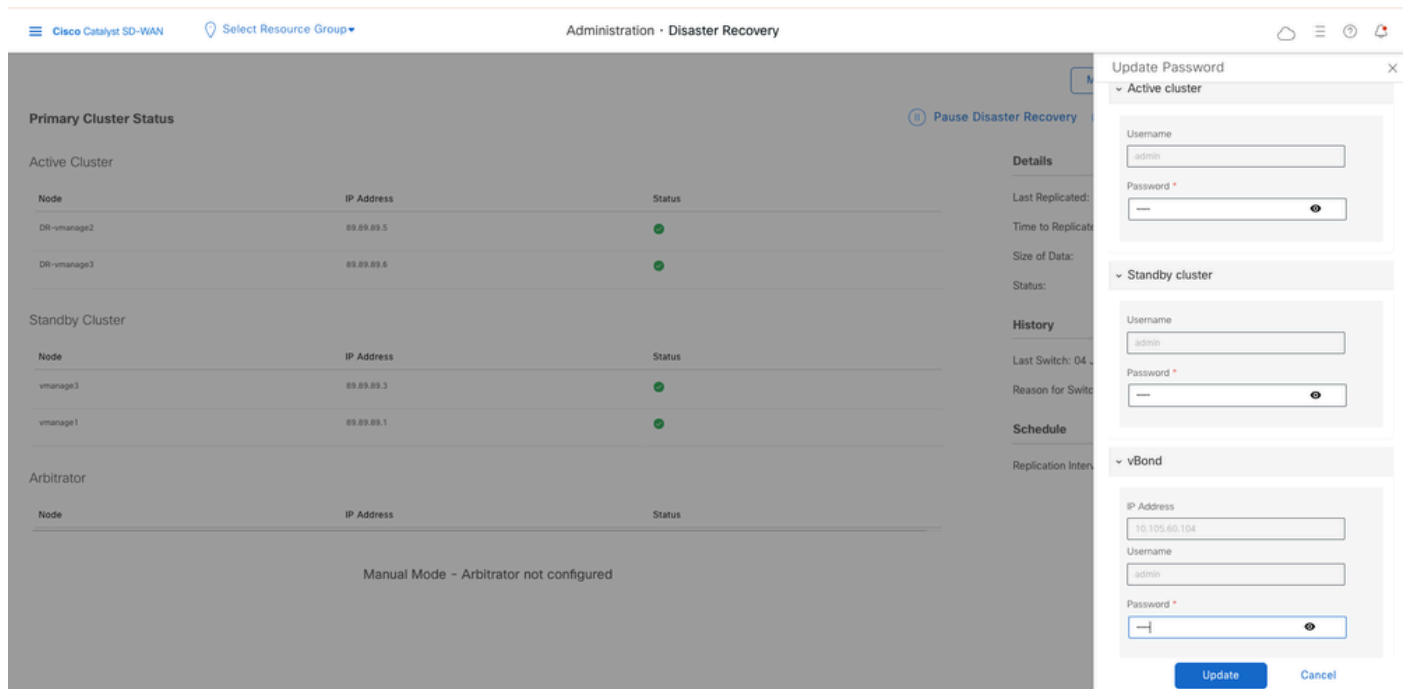
Update Validator (vbond) Password

Ensure to update the new vBond password on both the Disaster Recovery page and under Manage Password:

Administration > Disaster Recovery > Manage Password > Update vBond password.

Ensure replication is successful after updating the password. Attempt a failover only after confirming successful replication.

caveat: <https://bst.cloudapps.cisco.com/bugsearch/bug/CSCwn19224>.



Adding New Validator (vBond) to Overlay After Disaster Recovery Registration

Adding a new validator to the SD-WAN overlay after disaster recovery registration is not supported, as the disaster recovery setup is not aware of this new validator information since it was not updated during registration.

Although you can add the validator, a switchover fails.

If you need to add a new validator, observe these steps:

1. Delete the disaster recovery setup.
2. Add the new validator to the SD-WAN overlay.
3. Reconfigure disaster recovery.

Upgrade Disaster Recovery Overlays

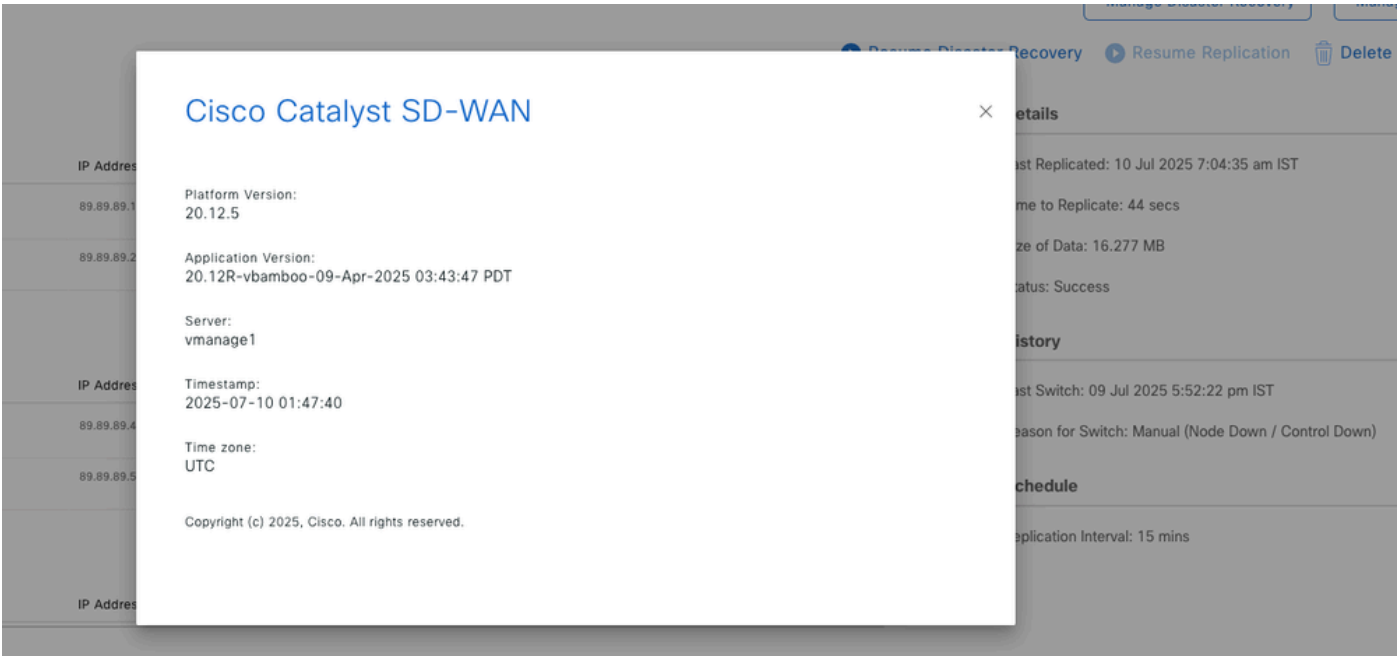
Before You Begin

- Use the CLI method to upgrade both the active and standby Cisco SD-WAN Managers.
- Ensure that the replication status on the **Administration > Disaster Recovery** page is stable and not in a transient state such as Import Pending, Export Pending, or Download Pending. It must be in the Success state before pausing disaster recovery.
- Pause the disaster recovery using **Pause Disaster Recovery** under **Administration > Disaster Recovery** page.

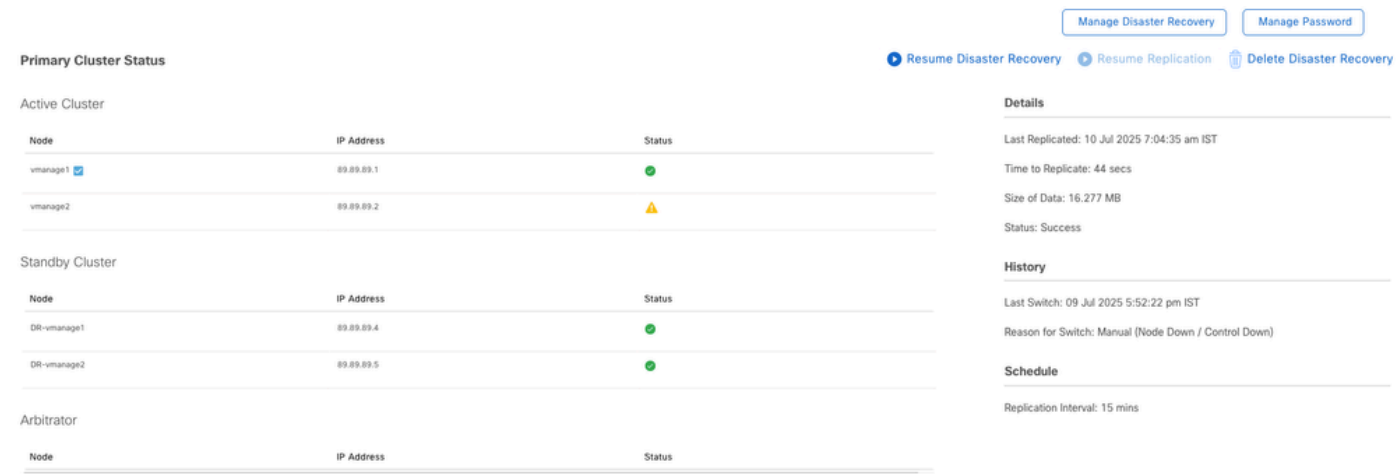
Upgrade Process

In this case you are upgrading vManage cluster from 20.12.5 to 20.15.2. Use the CLI method in order to upgrade the cluster.

Before you upgrade, verify version and replication status.



Pause disaster recovery:



After the upgrade, ensure all services are running and that you can log in to all vManage nodes (DC and DR) using the GUI.

Manage Disaster Recovery

Manage Password

Resume Disaster Recovery

Delete Disaster Recovery

IP Address

89.89.89.2

89.89.89.3

89.89.89.1

89.89.89.5

89.89.89.6

89.89.89.4

Cisco Catalyst SD-WAN

Platform Version:
20.15.2

Application Version:
20.15R-vbamboo-05-Mar-2025 01:53:17 PST

Server:
vmanage1

Timestamp:
2025-07-10 02:40:05

Time zone:
UTC

Copyright (c) 2025, Cisco. All rights reserved.

Close

Details

Last Replicated: 10 Jul 2025 7:04:35 AM GMT+5

Time to Replicate: 44 secs

Size of Data: 16.277 MB

Status: Success

History

Last Switch: 09 Jul 2025 5:52:22 PM GMT+05:30

Reason for Switch: Manual (Node Down / Control Down)

Schedule

Replication Interval: 15 mins

Resume disaster recovery; replication starts, and the replication status must eventually show as **success**.

The network is out of compliance due to licensing, please [click here](#) for more actions.

Snooze

Disaster Recovery

Manage Disaster Recovery

Manage Password

Pause Disaster Recovery

Delete Disaster Recovery

Primary Cluster Status

Active Cluster (3)

Node	IP Address	Status
vmanage2	89.89.89.2	✓
vmanage3	89.89.89.3	✓
vmanage1	89.89.89.1	✓

Standby Cluster (3)

Node	IP Address	Status
DR-vmanage2	89.89.89.5	✓
DR-vmanage3	89.89.89.6	✓
DR-vmanage1	89.89.89.4	✓

Details

Last Replicated: 10 Jul 2025 8:32:37 AM GMT+5

Time to Replicate: 46 secs

Size of Data: 16.401 MB

Status: Success

History

Last Switch: 09 Jul 2025 5:52:22 PM GMT+05:30

Reason for Switch: Manual (Node Down / Control Down)

Schedule

Replication Interval: 15 mins

Related Information

- <https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/ha-scaling/ios-xe-17/high-availability-book-xe/m-disaster-recovery.html>
- [Cisco Technical Support & Downloads](#)